

Hierarchical Shape Abstraction for Analysis of Free-List Memory Allocators

Bin Fang^{2,1} and Mihaela Sighireanu¹

¹ IRIF, University Paris Diderot and CNRS, France

² Shanghai Key Laboratory of Trustworthy Computing, ECNU, China

Abstract. We propose a hierarchical abstract domain for the analysis of free-list memory allocators that tracks shape and numerical properties about both the heap and the free lists. Our domain is based on Separation Logic extended with predicates that capture the pointer arithmetics constraints for the heap-list and the shape of the free-list. These predicates are combined using a hierarchical composition operator to specify the overlapping of the heap-list by the free-list. In addition to expressiveness, this operator leads to a compositional and compact representation of abstract values and simplifies the implementation of the abstract domain. The shape constraints are combined with numerical constraints over integer arrays to track properties about the allocation policies (best-fit, first-fit, etc). Such properties are out of the scope of the existing analyzers. We implemented this domain and we show its effectiveness on several implementations of free-list allocators.

1 Introduction

A dynamic memory allocator (DMA) is a piece of software managing a reserved region of the heap. It appears in general purpose libraries (e.g., C standard library) or as part of applications where the dynamic allocation shall be controlled to avoid failure due to memory exhaustion (e.g., embedded critical software). A client program interacts with the DMA by requesting blocks of memory of variable size that it may free at any time. To offer this service, the DMA manages the reserved memory region by partitioning it into arbitrary sized blocks of memory, also called *chunks*. When a chunk is allocated to a client program, the DMA can not relocate it to compact the memory region (like in garbage collectors) and it is unaware about the kind (type or value) of data stored. The set of chunks not in use, also called *free chunks*, is managed using different techniques. In this paper, we focus on *free-list allocators* [19, 26], that records free chunk in a list. This class of DMA includes textbook examples [17, 19] and real-world allocators [20].

The automated analysis of DMA faces several challenges. Although the code of DMA is not long (between one hundred to a thousand LOC), it is highly optimised to provide good performance. Low-level code (e.g., pointer arithmetics, bit fields, calls to system routines like `sbrk`) is used to manage efficiently (i.e., with low additional cost in memory and time) the operations on the chunks in the reserved memory region. At the same time, the free-list is manipulated

using high level operations over typed memory blocks (values of C structures) by mutating pointer fields without pointer arithmetic. The analyser has to deal efficiently with this *polar usage of the heap* made by the DMA. The invariants maintained by the DMA are complex. The memory region is organised into a *heap-list* based on the size information stored in the chunk header such that chunk overlapping and memory leaks are avoided. The start addresses of chunks shall be aligned to the some given constant. The free-list may have complex shapes (cyclic, acyclic, doubly-linked) and may be sorted by the start address of chunks to ease free chunks coalescing at deallocation. A precise analysis shall keep track of both numerical and shape properties to infer specifications implying such invariants for the allocation and deallocation methods of the DMA.

These challenges have been addressed partially by several works in the last ten years [5, 23, 25]. Efficient numerical analyses has been designed to track address alignment and bit-fields [23]. The most important progress has been done by the analysis proposed by Calcagno et al [5]. It is able to track the free-list shape and the numerical properties of chunk start addresses due to an abstract domain built on an extension of Separation Logic (SL) [24] with numerical constraints and predicates specifying memory blocks. However, some properties of the heap and free-list can not be tracked, e.g., the absence of memory leaks or ordering of start addresses of free-chunks. Although the analysis in [25] does not concern DMA, it is the first to propose a hierarchical abstraction of the memory to track properties of linked data structures stored in static memory regions. However, this analysis can not track properties like address sorting of the high level data structures (here the free-list) stored in the memory region. Furthermore, its link with a logic theory is not clear. Thus, a precise, logic based analysis for the inference of properties of free-list DMA is still a challenge.

In this paper, we propose a static analysis that is able to infer the above complex invariants of DMA on both heap-list and free-list. We define an abstract domain which uses logic formulas to abstract DMA configurations. The logic proposed extends the fragment of symbolic heaps of SL with a hierarchical composition operator \ni to specify that the free-list cover partially the heap-list. This operator provides a hierarchical abstraction of the memory region under the DMA control: the low-level memory manipulations are specified at the level of the heap-list and propagated in a way controlled by the abstraction at the level of the free-list. The shape specification is combined with a fragment of an array theory to capture properties of chunks in lists, similar to [3]. This combination is done in an accurate way wrt logic by including sequences of chunk addresses in the inductive definitions of list segments. The main advantages and contributions of this work are (1) *high precision of the abstraction* which is able to capture complex properties of free-list DMA implementations, (2) *strong logical basis* allowing to infer invariants that may be used by other verification methods, and (3) *modularity* of the abstract domain permitting to reuse existing abstract domains for the analysis of linked lists with integer data.

2 Overview

Fig. 1 includes excerpts from our running example, a free-list DMA implementation proposed in [1]. The type `HDR` (Fig. 1 (a)) defines the informations stored by the DMA at the start of chunks. The `size` field stores the full size of the chunk (in blocks of `sizeof(HDR)` bytes) and is used by the heap-list to determine the start of the next chunk. The `fnx` field is valid only for free chunks (i.e., chunks in the free-list) and stores the start address of the next free chunk. To simplify the presentation, we added the ghost field `isfree`, to mark explicitly free chunks. The memory region managed by the DMA is between addresses stored by the global variables `_hsta` and `_hend`; they are initialised by `minit` using `sbrk` calls. The start of the free list is stored in `frhd`. An intuitive view of the concrete state of the DMA is shown in Fig. 1(d). The busy chunks are represented in grey. The “next chunk” relation in the heap-list defined by the `size` field is represented by the lower arrows; the upper arrows represent the “next free chunk” relation defined by the `fnx` field. Furthermore, other structural invariants should be maintained after each call of DMA methods: the heap-list shall be correct inside the memory region `[_hsta, _hend)`, consecutive chunks of the heap-list are not both free (*early coalescing* policy), the free-list shall include only chunks of the heap-list, be acyclic and sorted by start address of chunks. The allocation method searches a chunk with size bigger than the requested `nbytes`; if the chunk is larger, it is split in two parts such that the last part (the end of the initial chunk) is allocated.

The goal of our analysis is to establish that, if the client uses correctly the DMA methods, these methods (i) preserve the above structural invariants and (ii) are memory safe. In particular, we analyse the DMA methods starting from a client program which initialises the DMA and then calls allocation and deallocation methods (see Sec. 5) in a correct way.

Heap-list abstraction. The concrete memory configurations managed by the DMA are abstracted by a first abstraction layer based on the *symbolic heap graphs* fragment of SL [9]. The fragment is parameterised by a set of predicates which capture the heap-list as follows:

- The predicate `blk(X; Y)`, introduced in [5], specifies an untyped sequence of bytes between the symbolic addresses X and Y . E.g., the configuration obtained at line 20 of `minit` is abstracted by `blk(_hsta; _hend)`.
- The predicate `chd(X; Y)` specifies a memory block `blk(X; Y)` storing a value of type `HDR`; the fields of this value are represented by the symbolic variables $X.size$, $X.fnx$, and $X.isfree$ respectively.
- The predicate `chk(X; Y)` specifies a chunk built from a chunk header `chd(X; Z)` followed by a block `blk(Z; Y)` such that the full memory occupied has size $Y - X = X.size \times \text{sizeof}(HDR)$.
- A well formed heap-list segment starting at address X and ending before Y is specified using the predicate `hls(X; Y)[W]`. The inductive definition of this predicate (see Tab. 2). requires that chunks does not overlap or leave memory leaks. The variable W registers the sequence of start addresses of


```

1 typedef struct hdr_s {
2     struct hdr_s *fnx;
3     size_t size;
4     //@ghost bool isfree;
5 } HDR;
6
7 static void *_hsta = NULL;
8 static void *_hend = NULL;
9 static HDR *frhd = NULL;
10 static size_t memleft;
11
12 void minit(size_t sz)
13 {
14     size_t align_sz;
15     align_sz = (sz+sizeof(HDR)-1)
16               & ~(sizeof(HDR)-1);
17
18     _hsta = sbrk(align_sz);
19     _hend = sbrk(0);
20
21     frhd = _hsta;
22     frhd->size = align_sz / sizeof(HDR);
23     frhd->fnx = NULL;
24     //@ghost frhd->isfree = true;
25
26     memleft = frhd->size;
27 }
28
29 void* malloc(size_t nbytes)
30 {
31     HDR *nxt, *prv;
32     size_t nunits =
33         (nbytes+sizeof(HDR)-1)/sizeof(HDR) + 1;
34
35     for (prv = NULL, nxt = frhd; nxt;
36          prv = nxt, nxt = nxt->fnx) {
37         if (nxt->size >= nunits) {
38             if (nxt->size > nunits) {
39                 nxt->size -= nunits;
40                 nxt += nxt->size;
41                 nxt->size = nunits;
42             } else {
43                 if (prv == NULL)
44                     frhd = nxt->fnx;
45                 else
46                     prv->fnx = nxt->fnx;
47             }
48             memleft -= nunits;
49             //@ghost nxt->isfree = false;
50             return ((void*)(nxt + 1));
51         }
52     }
53     warning("Allocation Failed!");
54     return (NULL);
55 }

```

(a) Globals and initialisation

(b) Allocation



(c) Part of the abstract invariant at line 34

(d) Concrete memory

Fig. 1. Running example with code, concrete memory, and abstract specification

chunks in the list segment and it is used to put additional constraints on the fields of these chunks. For DMA with early coalescing of free-chunks, we abstract the heap-list segments by a stronger predicate, *hlsc*.

These predicates are combined using the *separation conjunction operator* $*$ of SL, which requires disjointness of memory regions specified by its operands. The bottom of Fig. 1(c) illustrates the heap-list abstraction of the concrete memory provided in Fig. 1(d); for readability, the abstraction is represented by its Gaifman graph. The ghost variable *hli* represents the end of the data segment of the DMA, as returned by *sbrk*(0).

Hierarchical abstraction of the free-list. The first abstraction layer captures the total order of chunks in the heap-list. The free-list defines a total order over the set of free chunks. The second layer captures this order using the same SL fragment but over a different set of predicates (see Tab. 2):

- The predicate $\text{fck}(X; Y)$ specifies a chunk $\text{chk}(X; \dots)$ starting at X , with $X.\text{fnx}$ bound to Y and $X.\text{isfree}$ set to true.
- The predicate $\text{fls}(X; Y)[W]$ specifies a free-list segment starting at X , whose last element field fnx points to Y ; W registers the sequence of start addresses of free chunks in the list segment. The predicate $\text{flso}(X, \dots)[W]$ abstracts free-list segments sorted by the start address of chunks.

The top of Fig. 1(c) illustrates the free-list abstraction by its Gaifman graph.

Finally, the memory shape abstraction is obtained by composing the two abstraction levels using a new operator \ni , which requires that the set of chunks in the free-list abstraction is exactly the sub-set of chunks of the heap list with the isfree field set to true. Notice that the \ni operator can not be replaced by the logical conjunction because we are using the intuitive semantics of SL where spatial formulas fully specify the memory configurations. Or the free-list abstraction provides only a partial specification of the heap.

Constraints over sequences of chunk addresses. The inductive invariants specify invariants of DMA independent of user input, e.g., early chunk coalescing or free-list sorting. To capture policies like first-fit allocation (implemented by the `malloc` in Fig. 1(b)), we introduce universal constraints over sequences of chunk start addresses W attached to shape atoms, like in [3]. For example, the first-fit policy obtained at line 37 of `malloc`, is specified by:

$$\begin{aligned} \text{hlsc}(X_0; \text{hli})[W_H] \ni & (\text{fls}(Y_0; Y_2)[W_1] * \text{fck}(Y_2; Y_3) * \text{fls}(Y_3; \text{nil})[W_2]) \\ & \wedge Y_2.\text{size} \geq \text{nunits} \wedge \forall X \in W_1. X.\text{size} < \text{nunits} \end{aligned} \quad (1)$$

where Y_2 is the symbolic address stored in the program variable `nxt`. The general form of universal constraints is $\forall X \in W. A_G \Rightarrow A_U$ where A_G and A_U are arithmetic constraints over X and its fields. To obtain an efficient analysis, we fix A_G and infer A_U . We require that both A_G and A_U belong to a class of constraints allowed by a numerical abstract domain (see Sec. 3).

Static analysis with hierarchical shape abstraction. Overall, the analysis algorithm is a standard shape analysis algorithm. To expose fields constrained or assigned by the program statements, it unfolds predicate definitions. To limit the size of the abstraction, the algorithm normalises formulas to maintain only symbolic addresses that are cut-points, i.e., they are stored in the program variables or are sharing points in lists. This transformation of formulas folds back sub-formulas into more general predicates. Because the set of normalised shape formulas is bounded, we define widening only for the sequence constraints.

The hierarchical shape requires to solve a number of specific issues (see Sec. 5). The unfolding of shape predicates shall be done at the appropriate level of abstraction. For example, a traversal of the free-list requires only unfolding and folding at the free-list level. The heap-list level may abstract chunks and symbolic addresses which are in the free-list level. Thus, we define protocols for the unfolding and folding operations at each level that are sound wrt the hierarchical composition \ni and with the sequence constraints.

Table 1. Logic syntax

$X, Y \in \mathbf{AVar}$ location variables	$W \in \mathbf{SVar}$ sequence variables
$i, j \in \mathbf{IVar}$ integer variables	$\# \in \{=, \neq, \leq, \geq\}$ comparison operators
$x \in \mathbf{Var}$ logic variable	$\vec{x}, \vec{y} \in \mathbf{Var}^*$ vector of variables
$X.f$ field access	t, Δ integer term resp. formula
<hr/>	
$\varphi ::= \Pi \wedge \Sigma \mid \varphi \vee \varphi \mid \exists x. \varphi$	formulas
<hr/>	
$\Pi ::= A \mid \forall X \in W. A \Rightarrow A \mid W = w \mid \Pi \wedge \Pi$	pure formulas
$A ::= X[\mathbf{f}\mathbf{n}\mathbf{x}] - Y[\mathbf{f}\mathbf{n}\mathbf{x}] \# t \mid \Delta \mid A \wedge A$	location and integer constraints
$w ::= \epsilon \mid [X] \mid W \mid w.w$	sequence terms
<hr/>	
$\Sigma ::= \Sigma_H \ni \Sigma_F$	spatial formulas
$\Sigma_H ::= \mathbf{emp} \mid \mathbf{blk}(X; Y) \mid \mathbf{chd}(X; Y) \mid \mathbf{chk}(X; Y) \mid X \mapsto x \mid$ $\mathbf{hls}(X; Y)[W] \mid \mathbf{hlsc}(X, i; Y, j)[W] \mid \Sigma_H * \Sigma_H$	heap formulas
$\Sigma_F ::= \mathbf{emp} \mid \mathbf{fck}(X; Y) \mid \mathbf{fls}(X; Y)[W] \mid \mathbf{fso}(X, x; Y, y)[W] \mid \Sigma_F * \Sigma_F$	free-list formulas

3 Logic Fragment Underlying the Abstract Domain

We formalise in this section a fragment of Separation Logic [24] used to define the values of our abstract domain in Sec. 4.

Syntax. Let \mathbf{AVar} be a set of *location variables* representing heap addresses; to simplify the presentation, we consider that \mathbf{AVar} contains a special variable \mathbf{nil} representing the null address, also denoted by \mathbf{nil} . Let \mathbf{SVar} be a set of *sequence variables*, interpreted as sequences of heap addresses and \mathbf{IVar} be a set of *integer variables*. The full set of *logic variables* is denoted by $\mathbf{Var} = \mathbf{AVar} \cup \mathbf{SVar} \cup \mathbf{IVar}$. The domain of heap addresses is denoted by \mathbb{A} while the domain values stored in the heap is generically denoted by \mathbb{V} , thus $\mathbb{A} \subseteq \mathbb{V}$. To simplify the presentation, we fix \mathbf{HDR} , the type of chunk headers, and its fields $\{\mathbf{size}, \mathbf{fnx}, \mathbf{isfree}\}$ typed as declared in Fig. 1. The syntax of formulas is given in Tab. 1.

Formulas are in disjunctive normal form. Each disjunct is built from a pure formula Π and a spatial formula Σ . Pure formulas Π characterise the values of logic variables using comparisons between location variables, e.g., $X - Y = 0$, constraints Δ over integer terms, and sequence constraints. We let constraints in Δ unspecified, though we assume that they belong to decidable theories, e.g., linear arithmetic. The integer terms t are built over integer variables and field accesses using classic arithmetic operations and constants. We denote by Π_{\forall} (resp. Π_W, Π_{\exists}) the set of sub-formulas of Π built from universal constraints (resp. sequence constraints, quantifier free arithmetic constraints).

A spatial formula has two components: Σ_H specifies the heap-list and the locations outside this region; Σ_F specifies only the free-list. The operator \ni ensures that all locations specified by Σ_F are start addresses of free chunks in

Table 2. Derived predicates

$\text{chd}(X; Y) \triangleq \text{blk}(X; Y) \wedge \text{sizeof}(\text{HDR}) = Y - X \wedge X \equiv_{\text{sizeof}(\text{HDR})} 0$
$\text{chk}(X; Y) \triangleq \exists Z \cdot \text{chd}(X; Z) * \text{blk}(Z; Y) \wedge X.\text{size} \times \text{sizeof}(\text{HDR}) = Y - X$
$\text{fck}(X; Y) \triangleq \exists Z \cdot \text{chk}(X; Z) \wedge X.\text{isfree} = 1 \wedge X.\text{fnx} = Y$
$\text{hls}(X; Y)[W] \triangleq \text{emp} \wedge X = Y \wedge W = \epsilon$ $\vee \exists Z, W' \cdot \text{chk}(X; Z) * \text{hls}(Z; Y)[W'] \wedge W = [X].W'$
$\text{hlsc}(X, f_p; Y, f_\ell)[W] \triangleq \text{emp} \wedge X = Y \wedge W = \epsilon \wedge 0 \leq f_p + f_\ell \leq 1$ $\vee \exists Z, W', f \cdot \text{chk}(X; Z) * \text{hlsc}(Z, f; Y, f_\ell)[W'] \wedge W = [X].W'$ $\wedge f = X.\text{isfree} \wedge 0 \leq X.\text{isfree} + f_p \leq 1$
$\text{fls}(X; Y)[W] \triangleq \text{emp} \wedge X = Y \wedge W = \epsilon$ $\vee \exists Z, W' \cdot \text{fck}(X; Z) * \text{fls}(Z; Y)[W'] \wedge W = [X].W' \wedge X \neq Y$
$\text{fiso}(X, x; Y, y)[W] \triangleq \text{emp} \wedge X = Y \wedge W = \epsilon \wedge x - y \leq 0$ $\vee \exists Z, W' \cdot \text{fck}(X; Z) * \text{fiso}(Z, X; Y, y)[W'] \wedge W = [X].W' \wedge x - X \leq 0$

the heap-list. The atom emp holds iff the domain of the heap is empty. The *points-to atom* $X \mapsto x$ specifies that the heap contains exactly memory block at location X storing the value given by x . The *block atom* $\text{blk}(X; Y)$ holds iff the heap contains a block of memory at location X ending before the location Y . The other predicates are derived from blk and defined in Tab. 2. Notice that the *chunk header atom* $\text{chd}(X; Y)$ does not expose the fields of the block at location X using the points-to operator of SL. This ease the manipulation of heap-list level formulas, e.g., the coalescing of block and chunk atoms into a single block.

Semantics. Formulas φ are interpreted over pairs (I, h) where I is an *interpretation* of logic variables and h is a *heap* mapping a location to a non-empty sequence of values stored at this location. Formally, $I \in [(\text{AVar} \cup \text{IVar}) \rightarrow \mathbb{V}] \cup [\text{SVar} \rightarrow \mathbb{V}^*]$ and $h \in [\mathbb{A} \rightarrow \mathbb{V}^+]$ such that $\text{nil} \notin \text{dom}(h)$. Let $h(\ell)[i]$ denote the i th element of $h(\ell)$. Without loss of generality, we consider that a value of type HDR is a sequence of values indexed by fields. Tab. 3 provides the most important semantic rules. We denote by $\llbracket \varphi \rrbracket$ the pairs (I, h) such that $I, h \models \varphi$. The semantic entailment $\varphi \Rightarrow \psi$ is defined by $\llbracket \varphi \rrbracket \subseteq \llbracket \psi \rrbracket$.

Transformation rules. The definitions in Tab. 2 imply a set of lemmas used to transform formulas in abstract values (in Sec. 5). The first set of lemmas is obtained by directing predicate definitions in both directions. For example, each definition $P(\dots) \triangleq \vee_i \varphi_i$ introduces a set of *folding* lemmas $\varphi_i \Rightarrow P(\dots)$ and an *unfolding* lemma $P(\dots) \Rightarrow \vee_i \varphi_i$.

The second class of lemmas concerns list segment predicates in Tab. 2. The inductive definitions of these predicates satisfy the syntactic constraints defined in [12] for *compositional predicates*. Thus, every $P \in \{\text{hls}, \text{hlsc}, \text{fls}, \text{fiso}\}$ satisfies the following *segment composition lemma*:

$$P(X, \vec{x}; Y, \vec{y})[W_1] * P(Y, \vec{y}; Z, \vec{z})[W_2] \wedge W = W_1.W_2 \quad \Rightarrow \quad P(X, \vec{x}; Z, \vec{z})[W] \quad (2)$$

Table 3. Logic semantics: main rules

$I, h \models \Sigma_H \ni \Sigma_F$	iff	$I, h \models \Sigma_H$ and $\exists h' \subseteq h$ s.t. $I, h' \models \Sigma_F$ $\forall \ell \in \text{dom}(h') \cdot h'(\ell)[\text{isfree}] = 1$
$I, h \models \text{emp}$	iff	$\text{dom}(h) = \emptyset$
$I, h \models \text{blk}(X; Y)$	iff	$\text{dom}(h) = I(X) \wedge I(Y) - I(X) = h(I(X)) $
$I, h \models X \mapsto x$	iff	$\text{dom}(h) = I(X) \wedge h(I(X))[0] = I(x)$
$I, h \models \Sigma_1 * \Sigma_2$	iff	$\exists h_1, h_2$ s.t. $h = h_1 \uplus h_2$ and $I, h_i \models \Sigma_i$ for $i = 1, 2$
$I, h \models \forall X \in W \cdot A_1 \Rightarrow A_2$	iff	$I(W) = [a_1, \dots, a_n]$ s.t. $\forall i \in (1..n) \ I[X \mapsto a_i], h \models A_1 \Rightarrow A_2$
where		
$h_1 \subseteq h_2$	iff	$\text{dom}(h_1) \subseteq \text{dom}(h_2)$ and $\forall \ell \in \text{dom}(h_1) \cdot h_1(\ell) = h_2(\ell)$
$h_1 \otimes h_2$	iff	$\forall l_1 \in \text{dom}(h_1), l_2 \in \text{dom}(h_2) \cdot l_1 \neq l_2 \wedge$ $((l_1..l_1 + h_1(l_1) - 1) \cap (l_2..l_2 + h_2(l_2) - 1) = \emptyset)$
$h = h_1 \uplus h_2$	iff	$h_1 \otimes h_2, \text{dom}(h) = \text{dom}(h_1) \uplus \text{dom}(h_2)$, and $(h_1 \uplus h_2)(\ell) \triangleq \begin{cases} h_1(\ell) & \text{if } \ell \in \text{dom}(h_1) \\ h_2(\ell) & \text{if } \ell \in \text{dom}(h_2) \end{cases}$

The reverse implication is applied to split non-empty list segments. Finally, the block sub-formulas are removed, split, or folded using the following lemmas:

$$\text{blk}(X; Y) \wedge X \geq Y \Rightarrow \text{emp} \quad (3)$$

$$\text{blk}(X; Y) \wedge X < Y \Rightarrow \text{blk}(X; Z) * \text{blk}(Z; Y) \wedge X \leq Z \leq Y \quad (4)$$

$$\text{blk}(X; Y) * \text{blk}(Y'; Z) \wedge X \leq Y = Y' \leq Z \Rightarrow \text{blk}(X; Z). \quad (5)$$

4 Abstract Domain for Hierarchical Shape Abstraction

We define in this section the join-semilattice $\langle \mathcal{A}, \sqsubseteq, \sqcup \rangle$ used in our analysis. It is parameterised by a numerical abstract domain $\langle \mathcal{N}, \sqsubseteq^{\mathcal{N}}, \sqcup^{\mathcal{N}} \rangle$.

Concrete states. Let \mathbb{X} be the set of program variables, where **hli** is a ghost variable of location type. Values in \mathcal{A} represent sets of concrete states $M \in \mathbb{M}$ of the program. A concrete state M encloses an environment $\epsilon \in \mathbb{E} = \mathbb{X} \rightarrow \mathbb{A}$ mapping each program variable to its storing location, and a heap $h : \mathbb{A} \rightarrow \mathbb{V}^+$ mapping locations to sequences of values. For simplicity, **hli** symbol is overloaded to denote the symbolic location stored by **hli**.

Abstract values. Values in \mathcal{A} are a restricted form of logic formulas. Generally speaking, \mathcal{A} is a *co-fibered product* [6] of an *extended symbolic heap domain* for the spatial part and a *data word domain* [3] for the pure part. More precisely, \mathcal{A} includes a special value for \top and finite mappings of the form:

$$a^\# ::= \{ \langle \epsilon_i^\#, \Sigma_i(\vec{x}, \vec{W}) \rangle \mapsto \Pi_i(\vec{x}, \vec{W} \cup \{W_H, W_F\}) \}_{i \in I} \quad (6)$$

where $\epsilon_i^\# : \mathbb{X} \rightarrow \text{Var}$ is an abstract environment mapping program variables to symbolic location variables, Π_i includes arithmetic constraints allowed by \mathcal{N} , and the free variables of each formulas are detailed. Furthermore, the usage of sequence variables in Σ_i and Π_i are restricted as follows:

- R₁:** A sequence variable is bound to exactly one list segment atom in Σ_i ; thus Σ_i defines an injection between list segment atoms and sequence variables.
- R₂:** Π_i contains only the sequence constraints $W_H = w$ and $W_F = w'$, where W_H and W_F are special variables representing the full sequence of start addresses of chunks in the heap resp. free-list levels.

In addition, the universal constraints in the pure formulas Π_i are restricted such that, in any formula $\forall X \in W \cdot A_G \Rightarrow A_U$:

- R₃:** A_G and A_U use only terms where X appears inside a field access $X.f$.
- R₄:** A_G has one of the forms $X.size\#i$ or $X.isfree = i$.

These restrictions still permit to specify DMA policies like first-fit (see eq. (1)) and besides enable an efficient inference of universal constraints.

Internal representation. To ease the manipulation of extended spatial formulas $\langle \epsilon^\sharp, \Sigma \rangle$, we use the Gaifman graph representation, like in Fig. 1(c); nodes represent symbolic locations variables and labeled edges represent the spatial atoms in Σ or mappings in ϵ^\sharp . The universal formulas are represented by a map binding each pair of sequence variable and some guard A_G to a numerical abstract value.

Concretisation. An abstract value of the form (6) represents a formula $\forall_i \exists \vec{x}, \vec{W} \cdot \Sigma_i \wedge \Pi_i \wedge \epsilon_i^\sharp$ where each binding $(v, x) \in \epsilon_i^\sharp$ is encoded by $v \mapsto x$ (v is the location where is stored the program variable v). The false formula is represented by the empty mapping which corresponds to \perp . Based on the correspondence between abstract values and logic formulas, we define the concretisation $\gamma : \mathcal{A} \rightarrow \mathbb{M}$ by $\gamma(a^\sharp) = \llbracket a^\sharp \rrbracket$.

Ordering. The partial order \sqsubseteq is defined using a sound procedure inspired by [4, 12]. For any two non trivial abstract values $a^\sharp, b^\sharp \in \mathcal{A}$, $a^\sharp \sqsubseteq b^\sharp$ if for each binding $\langle \epsilon_i^\sharp, \Sigma_i \rangle \mapsto \Pi_i \in a^\sharp$ there exists a binding $\langle \epsilon_j^\sharp, \Sigma_j \rangle \mapsto \Pi_j \in b^\sharp$ such that:

- there is a graph isomorphism between the Gaifman graphs of spatial formula at each level of abstraction from Σ_i to Σ_j ; this isomorphism is defined by a bijection $\Psi : \text{img}(\epsilon_i^\sharp) \rightarrow \text{img}(\epsilon_j^\sharp)$ between symbolic location variables and a bijection Ω between sequence variables. Thus, $\Sigma_i[\Psi][\Omega] = \Sigma_j$,
- for each sequence constraint $W = w$ in $\Pi_{W,i}$, $\Omega(W) = \Omega(w)$ is a sequence constraint in $\Pi_{W,j}$,
- $\Psi(\Pi_{\exists,i}) \sqsubseteq^{\mathcal{N}} \Pi_{\exists,j}$,
- for each W defined in Σ_i and for each universal constraint $\forall X \in W \cdot A_G \Rightarrow A_U$ in $\Pi_{\forall,i}$, then $\Pi_{\forall,j}$ contains a universal constraint on $W' = \Omega(W)$ of the form $\forall X \in W' \cdot A_G \Rightarrow A'_U$ such that $\Psi(\Pi_{\exists,i} \wedge A_U) \sqsubseteq^{\mathcal{N}} A'_U$.

The following theorem is a consequence of restrictions on the syntax of formulas used in the abstract values.

Theorem 1 (\sqsubseteq soundness). *If $a^\sharp \sqsubseteq b^\sharp$ then $\gamma(a^\sharp) \subseteq \gamma(b^\sharp)$.*

Join. Given two non-trivial abstract values, a^\sharp and b^\sharp , their join is computed by joining the pure parts of bindings with isomorphic shape graphs [3]. Formally, for each two bindings $\langle \epsilon_i^\sharp, \Sigma_i \rangle \mapsto \Pi_i \in a^\sharp$ and $\langle \epsilon_j^\sharp, \Sigma_j \rangle \mapsto \Pi_j \in b^\sharp$ such that there is a graph isomorphism defined by Ψ and Ω between $\langle \epsilon_i^\sharp, \Sigma_i \rangle$ and $\langle \epsilon_j^\sharp, \Sigma_j \rangle$, we define their join to be the mapping $\{\langle \epsilon_j^\sharp, \Sigma_j \rangle \mapsto \Pi\}$ where Π is defined by:

- Π includes the sequence constraints of b^\sharp , i.e., $\Pi_W \triangleq \Pi_{W,j}$,
- $\Pi_\exists \triangleq \Psi(\Pi_{\exists,i}) \sqcup^\mathcal{N} \Pi_{\exists,j}$,
- for each W sequence variable in $\text{dom}(\Omega)$ and for each type of constraint A_G , then Π_\forall contains the formula $\forall X \in \Omega(W) \cdot A_G \Rightarrow \Psi(A_{U,i}) \sqcup^\mathcal{N} A_{U,j}$ where $A_{U,i}$ (resp. $A_{U,j}$) is the constraint bound to W (resp. $\Omega(W)$) in $\Pi_{\forall,i}$ (resp. $\Pi_{\forall,j}$) for guard A_G or \top if no such constraint exists.

The join of two bindings with non-isomorphic spatial parts is the union of the two bindings. Then, $(a^\sharp \sqcup b^\sharp)$ computes the join of bindings in a^\sharp with each binding in b^\sharp . Intuitively, the operator collects the disjuncts of a^\sharp and b^\sharp but replaces the disjuncts with isomorphic spatial parts by one disjunct which maps the spatial part to the join of the pure parts. Two universal constraints are joined when they concern the same sequence variables and guard A_G since $((\forall c \in W \cdot A_G \Rightarrow A_1) \vee (\forall c \in W \cdot A_G \Rightarrow A_2)) \Rightarrow (\forall c \in W \cdot A_G \Rightarrow (A_1 \vee A_2))$.

Theorem 2 (\sqcup soundness). *For any $a^\sharp, b^\sharp \in \mathcal{A}$, $\gamma(a^\sharp) \cup \gamma(b^\sharp) \subseteq \gamma(a^\sharp \sqcup b^\sharp)$.*

Cardinality of the abstract domain. The number of mappings in (6) increases during the symbolic execution by the introduction of new existential variables keeping track of the created chunks. Although the analysis stores only values with linear shape of lists (other shapes are signalled as error), the number of linear shapes is exponential in the number of nodes, in general. We avoid this memory explosion by eliminating existential variables using the transformation rules that replace sub-formulas by predicates, an operation classically called *predicate folding*. This operation uses lemmas (2)–(5), as discussed in Sec. 5. Thus, the domain of abstract values is bounded by an exponential on the number of pointer program variables local to DMA methods which is small in general, e.g., ≤ 3 in our benchmark. However, the domain of pure formulas used in the image of abstract values is not bounded because of integer constants. This fact requires to define widening operators for the data word domain used for the pure constraints.

5 Analysis Algorithm

We now describe the specific issues of the static analysis algorithm based on the hierarchical abstract domain presented in Sec. 4.

5.1 Main principles

The analysis is defined as a forward, non-relational abstract interpretation [8] over a shape abstract domain, and follows the principles of [7, 9, 10].

```

1 int main(void) {
2     minit(1024);
3     void* p = malloc(20);
4     malloc(20);
5     mfree(p); p = NULL;
6     p = malloc(20);
7     malloc(20);
8     mfree(p); p = NULL;
9     return 0;
10 }
    
```

Fig. 2. A client program

The DMA methods are analysed starting from some client function, like the main function in Fig. 2. The client programs are chosen to reveal policies of the DMA concerning the free chunk coalescing and the choice of the free chunk to be allocated. The analysis computes, for each control location ℓ , an abstract value $a_\ell^\#$ such that the $\gamma(a_\ell^\#)$ includes all the program configurations reachable by the concrete semantics of the program. For this, we define sound abstract transformers to compute abstract post-conditions and

widening operators to speed-up the convergence of the fix-point computation. The original points on abstract transformers concern the transfer of information between layers in the hierarchical unfolding, splitting, and folding of predicates, as detailed in Sec. 5.2–5.3. Furthermore, these operations are defined in a modular way, by extending [6] to data word numerical domains. The widening operator uses the widening of data word domain defined in [10].

5.2 Hierarchical unfolding

Abstract transfer functions compute over-approximations of post-images of atomic conditions and assignments in the program. For the spatial part, the abstract value is transformed such that the program variables read or written by the program operation are constrained using predicates that may capture the effect of the program operation. This transformation is called *predicate unfolding*.

We define the following partial order between predicates $\text{blk} < \text{chd} < \text{chk} < \text{fck} < \text{hls}, \text{hlsc}, \text{fls}, \text{flso}$ which intuitively corresponds to an increasing degree of specialisation. For each program operation s and each pointer variable x in s , an atom $P(X; \dots)$ with $\epsilon^\#(x) = X$ is transformed using lemmas in Sec. 3 to obtain the atom $Q(X; \dots)$ such that $Q \leq P$ is the maximal predicate satisfying:

- if s reads the fields of `HDR`, then $Q \leq \text{fck}$,
- if s assigns `x.isfree` or `x.fnx`, then $Q \leq \text{chk}$,
- if s mutates x using pointer arithmetics or assigns `x.size`, then $Q \leq \text{chd}$.

Consider the condition `nxt->size > nunits` at line 37 in Fig. 1(b), which reads the field `size`. Applied to the abstract value in Fig. 1(c), it requires to unfold the `flso` predicate from Y_2 , to obtain the top part of Fig. 3(a). To compute the post-image of the next assignment `nxt->size -= nunits`, the symbolic location Y_2 shall be the root a `chd` predicate (third case above). Thus, Y_2 is instantiated in the heap-list by first splitting and then unfolding the `hlsc` predicate, then by unfolding `chk` to obtain the bottom part of Fig. 3(a). The unfolding of `chk` requires to remove the `fck` atom from Y_2 in the free-list level because its definition is not more satisfied at the heap-list level.

The assignment `nxt += nxt->size` does not require to transform the predicate rooted in Y_2 because it is already $\leq \text{chd}$. Instead, the transformer adds a new symbolic location Z_1 in the heap-list level and constrain it by $Z_1 = Y_2 + Y_2.\text{size} \times \text{sizeof}(\text{HDR})$. If Z_1 goes beyond the limits of the user part

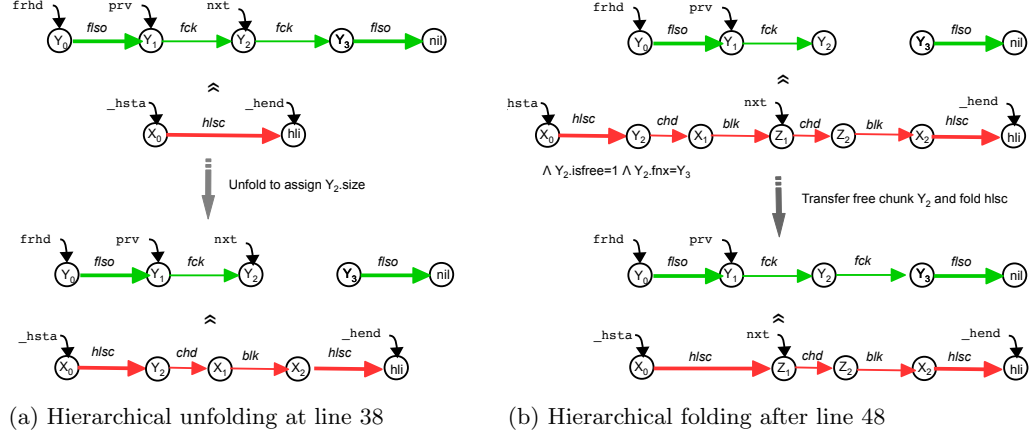


Fig. 3. Hierarchical abstract transformers for the running example

of the chunk starting at Y_2 (i.e., outside the interval $[X_1, X_2]$ in Fig. 3(a)), the analysis signals a chunk breaking. Otherwise, the `blk` atom from X_1 is split using lemma (4) to insert Z_1 ; the result is given by the top part of Fig. 3(b).

5.3 Hierarchical folding

To reduce the size of abstract values, the abstract transformers finish their computation on a binding $\langle \epsilon_i^\#, \Sigma_i \rangle \mapsto \Pi_i$ by eliminating the symbolic locations which are not cut-points in Σ_i . The elimination uses predicate folding lemma like (2) or (5) to replace sub-formulas using these variables by one predicate atom. The graph representation eases the computation of sub-formulas matching the left part of folding lemma.

More precisely, the elimination process has the following steps. First, it searches sequences of sub-formula of the form $\text{chd}(X_0; X_1) * \text{blk}(X_1; X_2) \dots * \text{blk}(X_{n-1}; X_n)$ where none of X_i ($1 \leq i < n$) is in $\text{img}(\epsilon^\#)$. Such sub-formulas are folded into $\text{chk}(X_0; X_n)$ if the pure part of the abstract value implies $X_0.\text{size} \times \text{sizeof}(\text{HDR}) = X_n - X_0$ (see Tab. 2). We use the variable elimination provided by the numerical domain \mathcal{N} to project out $\{X_1, \dots, X_{n-1}\}$ from the pure part. Furthermore, if the pure part implies $X_0.\text{isfree} = 1$, then the chunk atom (and its start address) is promoted as `fck` to the free-list level. This step is illustrated in Fig. 3(b) by sub-formulas $\text{chd}(Y_2; X_1) * \text{blk}(X_1; Z_1)$.

The second step folds list segments by applying their inductive definition and the composition lemma (2). The atoms $\text{chk}(X_0; \dots)$ for which the free-list level contains an atom $\text{fck}(X_0; \dots)$ may be folded at the heap-level into list segments due to the semantics of \exists . For example, the chunk starting from locations Y_2 is folded inside the heap-list segment in Fig. 3(b). Notice that folding of list segments implies the update of sequence and universal constraints like in [10].

Table 4. Benchmark of DMA

<i>DMA</i>	<i>LOC</i>	<i>List Pred.</i>	$ a^{\sharp} $	$ W_H / W_F $	<i>Invariants</i>
DKFF [19]	176	hlsc, flso	25	8/5	first-fit, MIN_SIZE-size
DKBF [19]	130	hlsc, flso	26	8/6	best-fit, MIN_SIZE-size
LA [1]	181	hlsc, flso	25	8/5	first-fit, 0-size
DKNF [19]	137	hlsc, flso	30	8/6	first-fit, MIN_SIZE-size
KR [17]	284	hlsc, flso	32	8/6	first-fit, 0-size

6 Experiments

We implemented the abstract domain in Ocaml and integrated it into an analyser as plug-in³ of the Frama-C platform [18]. We are using several modules of FRAMA-C, e.g., C parsing, abstract syntax tree transformations, and fix-point computation. The data word domain uses as numerical domain \mathcal{N} the library of polyhedra with congruence constraints provided by APRON [16]. To obtain precise numerical invariants, we transform program statements using bit-vector operations (e.g., line 16 of Fig. 1(a)) into statements permitted by the polyhedra domain that over-approximate the original effect.

We applied our analysis on the benchmark of free-list DMA in Tab. 4. DKFF and DKBF are implementations of Algorithms A and B from Sec. 2.5 of [19]. These DMA keep an acyclic free-list sorted by the start addresses of chunks. The deallocation does coalescing of successive free chunks. The allocation implements a first-fit resp. best-fit policy such that the fitting chunk is not split if the remaining free part is less than MIN_SIZE (variant proposed in [19]). This property is expressed by the invariant “MIN_SIZE-size” $\forall X \in W_H \cdot X.\text{size} \geq 8$ (here MIN_SIZE is 8 bytes) which is inferred by our analyser. The best-fit policy is implied by the universal constraint $\forall X \in W_i \cdot X.\text{size} \geq \text{rsz} \Rightarrow X.\text{size} \geq Y.\text{size}$ where rsz is the requested size, Y is the symbolic address of the fitting chunk, and W_i represents list segments around the fitting chunk. LA is our running example in Fig. 1; it follows the same principles as DKFF, but get rid of the constraint for chunk splitting. Our analyser infers the “0-size” invariant, i.e., $\forall X \in W_H \cdot X.\text{size} \geq 4$ ($=\text{sizeof}(\text{HDR})$). Notice that the code analysed fixes an obvious problem of the malloc method published in [1]. DKNF implements the next-fit policy using the “roving pointer” technique proposed in [19]: a global variable points to the chunk in the free-list involved in the last allocation or deallocation; malloc searches for a fitting free chunk starting from this pointer. Thus, the next-fit policy is a first-fit from the roving pointer. DKNF is challenging because the roving pointer introduces a case splitting that increases the size (number of disjuncts) in abstract values. The KR allocator [17] keeps a circular singly linked list, circularly sorted by the chunk start addresses; the start of the free-list points to the last deallocated block. The circular shape of the list requires to keep track of the free chunk with the biggest start address and this increases the size of the abstract values.

The analysis of each example with client program in Fig. 2 takes less than 3 seconds on a 2.53 GHz Intel Core 2 Duo laptop with 4GB of RAM. The

³ <https://www.irif.univ-paris-diderot.fr/~sighirea/celia/plus.html>

universally quantified invariants inferred for DMA policies are given in the last column. Columns $|a^\sharp|$ and $|W_H|/|W_F|$ provide the maximum number of disjuncts generated for an abstract value resp. the maximum number of predicate atoms in each abstraction level.

7 Related Works and Conclusion

Our analysis infers complex invariants of free-list DMA implementations due to the combination of two ingredients: the hierarchical representation of the shape of the memory region managed by the DMA and an abstract domain for the numerical constraints based on universally quantified formulas. The abstract domain has a clear logical definition, which facilitates the use of the inferred invariants by other verification methods.

The proposed abstract domain extends previous works [3, 5, 10, 11, 21]. We consider the SL fragment proposed in [5] to analyse programs using pointer arithmetic. We enrich this fragment in both spatial and pure formulas to infer a richer class of invariants. E.g., we add a heap-list level to track properties like chunk overlapping and universal constraints to infer first-fit policy invariants.

The split of shape abstraction on levels is inspired by works on overlaid data structures [11, 21]. We consider here a specific overlapping schema based on set inclusion which is adequate for DMA. We propose new abstract transformers which do not require user annotations like in [21]. Another hierarchical analysis of shape and numeric properties has been proposed in [25]. They consider the analysis of linked data structures coded in arrays and track the shape of these data structures and not the shape of the free set. Their approach is not based on logic and the invariants inferred on the content of list segments are simpler.

Our abstract domain includes a simpler version of the data word domain proposed in [3, 10], since the universal constraints quantify only one position in the list. Several abstract domains have been defined to infer invariants over arrays, e.g., [13] for array sizes, [14, 15] for array content. These works infer invariants of different kind on array partitions and they can not be applied directly to sequences of addresses. Recently, [22] defined an abstract domain for the analysis of array properties and applies it to the Minix 1.1 DMA, which uses chunks of fixed size. A modular combination of shape and numerical domains has been proposed in [6]. We extend their proposal to combine shape domains with domains on sequences of integers. Precise analyses exist for low level code in C [23] or for binary code [2]. They efficiently track properties about pointer alignment and memory region separations, but can not infer shape properties.

References

1. L. Aldridge. Memory allocation in C. *Embedded Systems Programming*, pages 35–42, August 2008.
2. G. Balakrishnan and T. W. Reps. Recency-abstraction for heap-allocated storage. In *SAS*, volume 4134 of *LNCS*, pages 221–239. Springer, 2006.
3. A. Bouajjani, C. Dragoi, C. Enea, and M. Sighireanu. On inter-procedural analysis of programs with lists and data. In *PLDI*, pages 578–589. ACM, 2011.

4. A. Bouajjani, C. Dragoi, C. Enea, and M. Sighireanu. Accurate invariant checking for programs manipulating lists and arrays with infinite data. In *ATVA*, volume 7561 of *LNCS*, pages 167–182. Springer, 2012.
5. C. Calcagno, D. Distefano, P. W. O’Hearn, and H. Yang. Beyond reachability: Shape abstraction in the presence of pointer arithmetic. In *SAS*, volume 4134 of *LNCS*, pages 182–203. Springer, 2006.
6. B. E. Chang and X. Rival. Modular construction of shape-numeric analyzers. In *Semantics, Abstract Interpretation, and Reasoning about Programs*, volume 129 of *EPTCS*, pages 161–185, 2013.
7. B. E. Chang, X. Rival, and G. C. Necula. Shape analysis with structural invariant checkers. In *SAS*, volume 4634 of *LNCS*, pages 384–401. Springer, 2007.
8. P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*, pages 238–252. ACM, 1977.
9. D. Distefano, P. W. O’Hearn, and H. Yang. A local shape analysis based on separation logic. In *TACAS*, volume 3920, pages 287–302. Springer, 2006.
10. C. Dragoi. *Automated verification of heap-manipulating programs with infinite data*. PhD thesis, University Paris Diderot, 2011.
11. C. Dragoi, C. Enea, and M. Sighireanu. Local shape analysis for overlaid data structures. In *SAS*, volume 7935 of *LNCS*, pages 150–171. Springer, 2013.
12. C. Enea, M. Sighireanu, and Z. Wu. On automated lemma generation for separation logic with inductive definitions. In *ATVA*, *LNCS*, pages 80–96. Springer, 2015.
13. S. Gulwani, T. Lev-Ami, and S. Sagiv. A combination framework for tracking partition sizes. In *POPL*, pages 239–251. ACM, 2009.
14. S. Gulwani, B. McCloskey, and A. Tiwari. Lifting abstract interpreters to quantified logical domains. In *POPL*, pages 235–246. ACM, 2008.
15. N. Halbwachs and M. Péron. Discovering properties about arrays in simple programs. In *PLDI*, pages 339–348. ACM, 2008.
16. B. Jeannet and A. Miné. Apron: A library of numerical abstract domains for static analysis. In *CAV*, volume 5643 of *LNCS*, pages 661–667. Springer, 2009.
17. B. W. Kernighan and D. Ritchie. *The C Programming Language, Second Edition*. Prentice-Hall, 1988.
18. F. Kirchner, N. Kosmatov, V. Prevosto, J. Signoles, and B. Yakobowski. Frama-C: A software analysis perspective. *FAC*, 27(3):573–609, 2015.
19. D. E. Knuth. *The Art of Computer Programming, Volume I: Fundamental Algorithms, 2nd Edition*. Addison-Wesley, 1973.
20. D. Lea. `dlmalloc`. <http://gee.cs.oswego.edu/pub/misc/malloc.c>, 2012.
21. O. Lee, H. Yang, and R. Petersen. Program analysis for overlaid data structures. In *CAV*, volume 6806 of *LNCS*, pages 592–608. Springer, 2011.
22. J. Liu and X. Rival. Abstraction of arrays based on non contiguous partitions. In *VMCAI*, volume 8931 of *LNCS*, pages 282–299. Springer, 2015.
23. A. Miné. Field-sensitive value analysis of embedded C programs with union types and pointer arithmetics. In *LCTES*, pages 54–63. ACM, 2006.
24. P. W. O’Hearn, J. C. Reynolds, and H. Yang. Local reasoning about programs that alter data structures. In *CSL*, *LNCS*, pages 1–19. Springer, 2001.
25. P. Sotin and X. Rival. Hierarchical shape abstraction of dynamic structures in static blocks. In *APLAS*, volume 7705 of *LNCS*, pages 131–147. Springer, 2012.
26. P. R. Wilson, M. S. Johnstone, M. Neely, and D. Boles. Dynamic storage allocation: A survey and critical review. In *IWMM*, volume 986 of *LNCS*, pages 1–116. Springer, 1995.